Subject: Re: [ckrm-tech] [RFC][PATCH 5/7] UBC: kernel memory accounting (core)
Posted by Chandra Seetharaman on Fri, 18 Aug 2006 20:26:19 GMT
View Forum Message <> Reply to Message

Kirill,

IMO, a UBC with resource constraint(limit in this case) should behave no
different than a kernel with limited memory. i.e it should do
reclamation before it starts failing allocation requests. It could even
do it preemptively.

There is no guarantee support which is required for providing QoS.

Each controller modifying the infrastructure code doesn't look good. We
can have proper interfaces to add a new resource controller.

chandra
On Wed, 2006-08-16 at 19:40 +0400, Kirill Korotaev wrote:
> Introduce UB_KMEMSIZE resource which accounts kernel
> objects allocated by task's request.
>
> Reference to UB is kept on struct page or slab object.
> For slabs each struct slab contains a set of pointers
> corresponding objects are charged to.
>
> Allocation charge rules:
> define1. Pages - if allocation is performed with __GFP_UBC flag - page
>   is charged to current's exec_ub.
> 2. Slabs - kmem_cache may be created with SLAB_UBC flag - in this
>   case each allocation is charged. Caches used by kmalloc are
>   created with SLAB_UBC | SLAB_UBC_NOCHARGE flags. In this case
>   only __GFP_UBC allocations are charged.
>
> Signed-Off-By: Pavel Emelianov <xemul@sw.ru>
> Signed-Off-By: Kirill Korotaev <dev@sw.ru>
>
<snip>
--

---------- ----------
Chandra Seetharaman | Be careful what you choose....
- sekharan@us.ibm.com | .....you may get it.
---------- ----------